

Royal Netherlands Academy of Arts and Sciences
P.O. Box 19121, 1000 GC Amsterdam, the Netherlands

Proceedings of the Colloquium
'Theory of Markets and Their Functioning'
Amsterdam, 1 - 3 July 1998

Supported by
Center for Economic Research, Tilburg
Tinbergen Research Institute, Amsterdam-Rotterdam

Koninklijke Nederlandse Akademie van Wetenschappen
Verhandelingen, Afd. Letterkunde, Nieuwe Reeks, deel 177

The Theory of Markets

Edited by P.J.J. Herings, G. van der Laan and A.J.J. Talman

AN EVOLUTIONARY MODEL OF MARKET STRUCTURE¹

CARLOS ALÓS-FERRER, ANA B. ANIA, FERNANDO VEGA-REDONDO

We study a market for a homogeneous good in which firms adjust their production decisions on the basis of imitation, learning from own experience, and local experimentation. For any fixed set of firms (more than one), long run behavior settles on a symmetric marginal cost pricing equilibrium. When market entry and exit are allowed, we find a sharp effect of technology on long-run market structure. Specifically, we show that, under decreasing returns and some fixed cost, the market grows to "full capacity" at Walrasian equilibrium; on the other hand, if returns are increasing, the unique long run outcome involves a profit-maximizing monopolist.

KEYWORDS: Local mutation, imitation dynamics, Cournot oligopoly, marginal cost pricing.

JEL CODES: C72, D43, L13.

1. INTRODUCTION

The traditional Theory of Industrial Organization has adopted the full rationality of firms as one of its basic premises. This has been widely criticized because it involves an excessive degree of rationality, extensive knowledge, and high computation capabilities. In contrast, evolutionary models view firms as agents whose rationality is bounded, due to some limited reasoning capacity or imperfect knowledge of the environment.

Alchian (1950) was the first to point out that imitation of success should be regarded as a major determinant of behavior in economic environments. For, among other reasons, it requires minimal knowledge on market conditions and imposes very little computational burden on firms. As a measure of success, he also argued, profits should be the key variable used by firms within a market environment.

A model with these characteristics has been recently proposed by one of us (Vega-Redondo (1997)), thereafter labelled VR.² This paper considers a market for a homogeneous good with n firms having access to the same technology. Every period, firms may imitate those outputs which led to the highest profits in the preceding period. Furthermore, they occasionally experiment (or "mutate") with

¹We thank Luis Corchón and Klaus Reiner Schenk-Hoppé for helpful comments. Financial support from the Instituto Valenciano de Investigaciones Económicas and from the Ministry of Education project no. PB94-1504 is gratefully acknowledged.

²This work borrows from recent evolutionary literature (see Kandori, Mailath and Rob (1993) or Young (1993)) some of its essential features, both conceptual and technical. For recent monographs on Stochastic Evolutionary Theory, the reader can refer to Samuelson (1997) or Vega-Redondo (1996).

some independent and small probability. Under the assumption that a symmetric Walrasian equilibrium exists in this market, VR shows that such an equilibrium is the unique stochastically stable state of the process (i.e. the only one visited a significant fraction of time in the long run). Thus, even in a market with a potentially small number of firms, the simple and intuitive behavioral rule "imitation of success" is seen to lead (when slightly perturbed) to a competitive outcome. This stands in sharp contrast with the conclusions obtained within the standard Cournot model under the assumption of full rationality on the part of firms.³

The present paper extends the analysis carried by VR in two important respects. First, we explore the implications of the evolutionary approach described⁴ to contexts where a Walrasian equilibrium does not exist; specifically, we characterize the long-run behavior arising in those situations where firms enjoy increasing returns throughout. Second, we introduce the possibility of population turnover (i.e. market entry and exit) and explore the implications of the underlying technological conditions on the long-run market structure.

Our results can be briefly summarized as follows. First, we find that, as long as there is more than one firm in the industry, the combination of imitation and experimentation leads firms to a symmetric state where everyone produces at a marginal cost which equals the market-clearing price. This, of course, reproduces the result of VR when a symmetric Walrasian equilibrium exists, since the equality of price and marginal cost is verified at such an equilibrium. However, when increasing returns prevail throughout (i.e. marginal cost is monotonically decreasing), this represents a substantial extension with new, rather surprising, implications. For example, it implies that incumbent firms will be forced into negative profits (i.e. losses) when increasing returns prevail.⁵

The second contribution of the present paper is to characterize the long-run market structure, as a function of the underlying technological conditions. We explore two general scenarios:

- (i) *decreasing returns*, i.e. decreasing marginal costs and a certain fixed cost;
- (ii) *increasing returns*, i.e. decreasing marginal costs.

³Similar considerations are discussed in Rhode and Stegeman (1994) or Schaffer (1989) within a restricted scenario with two firms and specific conditions on costs and demand. As explained in VR, the essential mechanism here involves certain considerations of "spite" arising in finite population evolutionary models.

⁴There are, however, two variations on the received evolutionary approach (in particular, that of VR) which are of some independent interest. First, we allow for arbitrarily long memory on the part of firms in adjusting their output. Second, we restrict experimentation to be local, i.e. to involve only "slight" deviations from the original output.

⁵The Marginal Cost Pricing Equilibrium has been proposed in General Equilibrium Theory as a normative rule to be used by firms under non-convexities in production. Its essential interest relies on the fact that one can show it to be a necessary (although not sufficient) condition for efficiency - see e.g. Quinzii (1992). Here, however, marginal cost pricing is obtained as a *positive* solution, so that our conclusions may be interpreted as providing some foundations for it.

We find a sharp ("knife-edge") effect of technology on market structure. In scenario (i), the process uniquely settles on the symmetric Walrasian equilibrium with very low or zero profits (i.e. the Walrasian outcome at full market capacity). In contrast, scenario (ii) is seen to induce a unique long-run outcome with a single monopolist in the market producing the output that maximizes profit.

These conclusions contrast with some of the "folk" ideas derived from the received Theory of Industrial Organization. There, one typically finds a *gradual* relationship between the degree of decreasing returns prevailing in the industry and the competitiveness of the induced outcome (i.e. its proximity to a Walrasian equilibrium). Consider for example, the well-known work of Novshek (1980). He shows that as the efficient scale of production falls relative to the size of the market (e.g. marginal costs become steeper), the number of active firms grow and the outcome approaches a Walrasian equilibrium. Conversely, when the efficient scale rises, fewer firms stay active in the market and, therefore, the less competitive gradually becomes the corresponding Cournot outcome.⁶

In a sense, our analysis reflects a view on market behavior which is reminiscent of that espoused by the pre-strategic Theory of Industrial Organization. If production returns are decreasing (and, therefore, price-taking behavior is well defined), a Walrasian outcome is obtained. In the opposite case where returns are increasing (no matter how mildly so), a "natural" and fully exploitative monopoly results which is immune to entry, actual or potential.⁷

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 carries out the analysis with a fixed number of firms. Section 4 augments the model to accommodate entry and exit and undertakes the corresponding analysis. Section 5 includes a general overview and discussion of the different results, suggesting as well some possible extensions. A summary of graph-theoretic techniques used in the analysis are summarized in Appendix 1. Finally, the formal proofs of the results are included in Appendix 2.

2. THE BASIC MODEL

2.1. A Market for a Homogeneous Good

Consider the market for a homogeneous product with $n \geq 2$ firms, $j \in N = \{1, \dots, n\}$. The demand side of the market will be modelled by an inverse demand

⁶Novshek's framework contemplates U-shaped average costs and is therefore incompatible with increasing marginal costs throughout, as considered in Subsection 4.1 below. However, this possibility could be readily introduced in his setup, leading to considerations analogous to those described above.

⁷This contrasts, for example, with the modern Theory of Contestable Markets (see, e.g. Baumol, Panzar and Willig (1982)).

function

$$P : \mathbb{R}_+ \rightarrow \mathbb{R}_+$$

which is assumed differentiable with $P'(\cdot) < 0$ and $P(x) \rightarrow 0$ as $x \rightarrow \infty$. All firms use the same technology to produce the good, as given by the cost function

$$C : \mathbb{R}_+ \rightarrow \mathbb{R}_+$$

which is taken to be twice differentiable and non-decreasing. Costs and demand are assumed to verify the following assumption.

$$A.1 \quad 2P'(nx) < C''(x) \quad \forall x.$$

This assumption⁸ requires that marginal costs do not decrease too rapidly relative to the demand function. It allows for increasing returns to scale (i.e. decreasing marginal cost) provided they are not too acute and, moreover, it is obviously satisfied under decreasing returns to scale (i.e., if $C''(x) \geq 0$ for all x). Standard Cournot oligopoly models (see for example Friedman (1982)) typically require for equilibrium existence that $P'(\sum_{k=1}^n x_k) - C''(x_i) < 0 \quad \forall x_1, \dots, x_n$, which trivially implies Assumption 1.

2.2. Firms' Dynamic Behavior

We postulate an evolutionary dynamics in discrete time $t = 0, 1, 2, \dots$ where, each period, firms produce some output level according to the given technology. For the sake of simplicity, we will assume that the output levels are chosen from a finite grid $\Gamma(\delta) = \{0, \delta, 2\delta, \dots, v\delta\}$ for some given $\delta > 0$, arbitrarily small, and some $v = v(\delta) \in \mathbb{N}$, arbitrarily large and such that $v\delta = K > 0$ fixed. This is a technical assumption motivated by our desire to remain within a simple framework with a finite number of possible states. One can think of δ as some indivisibility level or minimum production scale, and of K as the maximum relevant output which firms may ever consider.⁹

Next, we introduce firms' adjustment process, which reflects both considerations of imitation and occasional experimentation. Each of them is presented in turn.

⁸Note that it depends on the number of firms.

⁹No specific properties have been postulated on the inverse demand function that would bound the size of the market. However, this could be done naturally by assuming that the market-clearing price becomes negligible beyond some bounded interval and that no firm will ever consider producing a larger output.

Imitation Dynamics

Let $x_j(t)$ be the output level chosen by firm j at time t . Its profits at t are given by

$$\Pi_j(t) = P\left(\sum_{k=1}^n x_k(t)\right) x_j(t) - C(x_j(t)).$$

It will be postulated that when any firm is in the position to revise its output, it simply mimics one of those outputs which, given the information it has available, has produced the highest profit.¹⁰ More precisely, this imitation dynamics can be decomposed as follows:

- *Information*: At the end of each t , every firm is assumed to have information on the results (outputs and profits) of the last k periods (including the one just completed). However, we allow for the possibility that there might exist informational asymmetries, the most recent outputs and profits associated to any given firm being observed by the other firms only with a delay of s periods, $0 \leq s < k$. If $s = 0$, there is complete information, but if $s > 0$ then the most recent results of a firm are its private information.

Under this formulation, the profits known by firm j at the end of period t are those in the set

$$I_j(t) = Q(t) \cup R_j(t)$$

where

$$Q(t) = \{\Pi_i(t - \tau) : \tau = s, \dots, k - 1, i = 1, \dots, n\} \quad (1)$$

are the profits publicly known at the end of t , and

$$R_j(t) = \{\Pi_j(t - \tau) : \tau = 0, \dots, s - 1\} \quad (2)$$

are the (own) profits that remain private information for firm j .

- *Revision opportunities*: At the end of every period t , some non-empty subset of firms is given the option to revise their respective outputs for the next

¹⁰Imitation has been one of the most common behavioral paradigms used to model boundedly rational agents in recent Evolutionary Theory. See, for example, Björnerstedt and Schlag (1996), Robson and Vega-Redondo (1996), or Schlag (1998).

period $(t + 1)$. For simplicity, every such subset is selected with positive probability.¹¹

- *Adjustment*: If a firm is able to revise its output, it will simply mimic one of the outputs yielding highest profits among those it knows, i.e. one of the outputs in the set

$$B_j(t) = \{x_i(t - \tau) : \Pi_i(t - \tau) \in \text{Max } I_j(t)\}.$$

If this set is not a singleton, each of its elements are assumed chosen according to a firm-independent probability distribution with full support.

Experimentation

Each firm is assumed to occasionally experiment (tremble, mutate). Specifically, once the imitation-based adjustment has been completed, every firm is supposed subject to an independent probability $\epsilon > 0$ of changing its output "slightly", i.e. it increases or decreases output by an amount δ , both alternatives having positive probability in the interior of $\Gamma(\delta)$.¹²

3. ANALYSIS

We will denote

$$\Omega = \Gamma(\delta)^n \times \dots \times \Gamma(\delta)^n,$$

the state space of the dynamics.

¹¹We do not allow the empty subset to be selected in order to have always "meaningful" time periods: a period such that no firm is even given the opportunity to change its output will just be ignored.

This could be formally described as follows: Let $z(t)$ be an n -dimensional random variable taking values on the set $\{0, 1\}^n - \{(0, \dots, 0)\}$, and such that

$$\text{Pr}(z(t) = z) > 0 \quad \forall \quad z;$$

$z_j(t) = 1$ means that firm j has, in period t , the opportunity to revise its output for $t + 1$, while $z_j(t) = 0$ means that it is not possible to do so.

One example of such $z(t)$ would be to take the uniform distribution on $\{0, 1\}^n$, delete the state $(0, \dots, 0)$ and re-normalize the probabilities.

¹²This is a key difference with VR, where firms are allowed to mutate to any output in the grid. Here, however, we want to think of mutation as gradual experimentation. This idea could have also been formalized, for example, by postulating a probability $\epsilon^l > 0$ for changing output by $l \geq 1$ steps of size δ . In this way, the process would have been ergodic trivially and our results would also hold. The chosen approach has the advantage of reflecting an analogous idea of gradualness in a more clear-cut way.

At each t , the state of the system will be given by¹³

$$\begin{aligned} \omega(t) &= (\omega_1(t), \dots, \omega_k(t)) \\ &= [(x_1(t - k + 1), \dots, x_n(t - k + 1)), \dots, (x_1(t), \dots, x_n(t))]. \end{aligned}$$

States where each firm chooses the same output level during the k represented periods, i.e. $\omega_1(t) = \dots = \omega_k(t)$, will be called *repeated*. States where all firms choose the same output level during the k represented periods will be called *monomorphic*. The monomorphic state associated to the output y is denoted by $\omega(y) = [(y, \dots, y), \dots, (y, \dots, y)]$.

Although the transition probabilities depend not only on the last period but on the k last ones, the specification of the state space is such that all the relevant information for these transitions is contained in the previous state. So, the dynamics described defines a Markov Process with a finite state space, i.e. a Markov recurrent communication class which is aperiodic. Therefore, there exists a unique invariant distribution, associated to this recurrent class, which gives us the long run frequencies with which each state is observed, independently of initial conditions. To reflect its dependence on the experimentation rate and the density of the output grid, this invariant distribution is denoted by $\mu_{\epsilon, \delta}$.

Intuitively, we want to think of the experimentation probability ϵ as small. Moreover, we would like to make sure that the technical convenience afforded by a finite grid does not have any distorting effect on the analysis. With these two considerations in mind, the analysis proceeds as follows.

First, we fix a δ small enough and focus our analysis on the limit invariant distribution $\mu_\delta^* \equiv \lim_{\epsilon \rightarrow 0} \mu_{\epsilon, \delta}$, which is seen to be well defined. Those states in the support of μ_δ^* are called *stochastically stable states*. When the experimentation rate becomes small, it is only these states which are observed a significant fraction of time (a.s.) along any sample path of the process.

Secondly, we by-pass any artificial considerations which could be associated to the discreteness of the grid by focusing on the limit invariant distribution $\mu^* \equiv \lim_{\delta \rightarrow 0} \mu_\delta^*$, i.e. we consider an arbitrarily fine grid. We do this in a similar vein to the approach undertaken by Young (1993) in a different framework.

As established by Theorem 3.2 below, the whole mass of μ^* is concentrated on the monomorphic state where firms are at a Marginal Cost Pricing Equilibrium. Next, we state formally this key equilibrium concept.

DEFINITION 3.1: A (symmetric) *Marginal Cost Pricing Equilibrium* (MCPE) is a pair (y^*, p^*) such that $p^* = P(ny^*) = C'(y^*)$.

¹³Every state must include the description of the current and $k - 1$ preceding time periods because the output adjustment for $t + 1$ requires information from $t - k + 1$ to t .

If $C''(x) \geq 0$ for all x (i.e. costs are convex), then $y^* \in \arg \max_x [p^*x - C(x)]$ and (y^*, p^*) can be thought of as a *Walrasian Equilibrium*. On the other hand, if $C''(x) < 0$ for every x (i.e. under increasing returns), Walrasian equilibria do not exist, but a MCPPE still exists under quite general conditions. Rather than making them explicit, we simply adopt the following assumption:

A.2 There exists an MCPPE (y^*, p^*) with $y^* \in (0, K)$.

Note that, under A.1, $P(\pi x) - C'(x)$ is a strictly monotone function. Therefore, in view of A.2, we may conclude that, under our maintained assumption, the MCPPE is unique.

Our first result shows that the long-run prediction for this market is that all the firms produce the output level y^* . More precisely, in the proof of this result, we show that, with a fixed δ , the stochastically stable states are monomorphic states associated to output levels contained in a certain interval around y^* . This interval shrinks as δ becomes small, a property which can be seen as reflecting precision problems and that is formally captured by the fact that the limit when δ goes to zero of this interval is precisely the output level y^* .

THEOREM 3.2: *Assume A.1, A.2. Then, μ^* is well defined and $\mu^*(\varpi(y^*)) = 1$.*

PROOF: See Appendix 2.

Theorem 3.2 establishes that, in the long run, the market will spend "most" of its time in the monomorphic situation where all firms produce the same output level of the homogeneous good; moreover, this output is the one for which the market-clearing price equals marginal cost. In particular, with convex costs, this effectively reproduces the result of VR since, in a somewhat different context,¹⁴ it also selects the Walrasian Equilibrium. Nevertheless, it is striking to note that this conclusion (i.e. marginal-cost pricing) also holds under increasing returns to scale, a situation where it typically provides losses. It is precisely this observation that leads us to enrich our dynamical process with endogenous entry and exit of firms. As formulated next, in this augmented process firms will be allowed to exit from the market when losses are realized and, reciprocally, new firms will be allowed to enter when the market signals to outsiders that positive profits may be achieved in it.

4. ENTRY AND EXIT

Both the ideas of a firm sustaining losses without exiting the market, and a (non-regulated) market that shows positive profits without experiencing any entry of new firms seem unsatisfactory. To address these concerns, we propose the following enrichment of the original framework.

Fix an arbitrarily large maximum number of firms F .¹⁵ The new state space for the reformulated dynamics is

$$\hat{\Omega} = (\cup_{n=1}^F \Gamma(\delta)^n \cup \{\theta_0\}) \times \dots \times (\cup_{n=1}^F \Gamma(\delta)^n \cup \{\theta_0\})$$

where θ_0 is the trivial configuration where there are no firms in the market. On this space we consider the following extension of the formerly postulated dynamics.

First, at each t , those firms present in the market undertake processes of imitation and experimentation, as described in Subsection 2.2. Then, the set of firms participating in the market is modified through exit and entry as follows.

On the one hand, some incumbent firms are assumed to receive the option of revising their participation in the market. Formally, we postulate that every one of them has an independent probability $\sigma > 0$ of enjoying such revision opportunity. In that event, we simply assume that a firm decides to exit if, and only if, it is currently incurring losses.

Entry, on the other hand, is also formalized stochastically. Specifically, we suppose that, at the end of each t , there is some maximum number of potential firms ready to enter the market.¹⁶ If profits are observed in $Q(t)$ – recall (1), with the obvious re-interpretation – each potential entrant is subject to an independent probability $\eta > 0$ of effectively entering the market in $t + 1$. In that event, it is assumed to imitate one of the outputs in that set yielding maximum profits at t (i.e. displays a behavior equivalent to that of incumbents). If the process is at state $\varpi(\theta_0)$ – i.e. $Q(t) = \emptyset$ – we simply assume that such "market void" is filled by new firms, again entering with independent probability η . Since they have no market information to rely upon in shaping their starting decision, it is natural to assume that they choose some output from $\Gamma(\delta)$ according to a given probability distribution.

Since exit and entry introduce two additional sources of noise into the process (respectively associated to the probabilities σ and η above) their relative magnitude have to be specified. Intuitively, it seems reasonable to postulate that output adjustment by incumbent firms (either through imitation or experimentation) should be more flexible than population dynamics (i.e. entry or exit). In terms of the corresponding probabilities of occurrence of these events, this amounts to postulating

¹⁵This restriction is made to remain within a finite-state formulation. It could be endogenized, for example, by introducing natural conditions on demand (cf. Footnote 9).

¹⁶Again, this potential number of firms is bounded appropriately (as a function of the current state) so that the total number of firms in the market may never exceed F (recall Footnote 15).

¹⁴Recall Footnote 4.

that, viewed as infinitesimals, σ and η should be of larger order (i.e. converge faster to zero) than ϵ . Thus, if we denote by \sim the relationship "be infinitesimally as large" (i.e. an infinitesimal of no larger order), we shall assume:

$$\epsilon \sim \sigma; \quad \epsilon \sim \eta. \quad (3)$$

In fact, it will be convenient to conceive σ and η as functions of ϵ and assume that they are of finite order in ϵ , i.e.

$$\exists q, q' \in \mathbb{N} : \lim_{\epsilon \rightarrow 0} \frac{\epsilon^q}{\sigma(\epsilon)} = \lim_{\epsilon \rightarrow 0} \frac{\epsilon^{q'}}{\eta(\epsilon)} = 0. \quad (4)$$

As for the relative magnitudes of σ and η , the gist of our conclusions is unaffected by them. However, it will sharpen the analysis to assume:¹⁷

$$\sigma \succ \eta, \quad (5)$$

where \succ stands for the relationship "be infinitesimally larger (strictly)". This latter condition may be motivated by the idea that any decision to be taken by an incumbent that is already in the market (even if this decision is exit) is much more likely than any decision to be taken by some agent external to the market (i.e. entry).

As it was the case for the original context, the presently augmented process is seen to be ergodic for any $\epsilon > 0$ and small enough $\delta > 0$. Specifically, it induces a unique invariant distribution $\hat{\mu}_{\epsilon, \delta}$ which summarizes its long-run evolution. For considerations already explained, the analysis will be concerned with the limit invariant distribution

$$\mu^{**} \equiv \lim_{\delta \rightarrow 0} \lim_{\epsilon \rightarrow 0} \hat{\mu}_{\epsilon, \delta}$$

which captures long-run performance for an infinitesimal experimentation rate and an arbitrarily fine grid.

As advanced, the conclusions of the present augmented model are sharply affected by the qualitative nature of the underlying technology, i.e. whether it exhibits increasing or decreasing returns. Each of these two alternative scenarios is addressed in turn.

¹⁷This requirement is essentially a simplifying condition, which could be largely dispensed with in our analysis. For example, our analysis for the decreasing-returns scenario would be completely unaffected if all probabilities ϵ , σ , and η were infinitesimals of the same order. Instead, this would affect slightly the analysis for the alternative scenario with increasing returns in that, even though the monopoly state would still be stochastically stable, other states involving a "few" firms would also be so. (For example, if two firms both producing any output between the monopoly and the marginal-cost pricing outputs incur losses, all stochastically stable states involve at most two firms.)

4.1. Increasing Returns

For the sake of simplicity, we assume that there is a unique output which maximizes monopoly profits (i.e. the profits obtained by a single incumbent).

T.1 The monopoly profits function $P(\hat{y}) \cdot x - C(x)$ has a unique local maximum \hat{y} on $[0, K]$, with $P(\hat{y}) \cdot \hat{y} - C(\hat{y}) > 0$.

This is merely a technical assumption which does not affect the essence of our analysis.¹⁸ It is implied, for example, by the standard condition of strict concavity on the profit functions, often found in the Theory of Industrial Organization.

Let $\hat{\omega} \equiv (\hat{y}, \dots, \hat{y})$ be the repeated state where a monopoly chooses its profit-maximizing output. Under increasing returns, our conclusions are contained in the following result.

THEOREM 4.1: Suppose A.1 holds for all $n = 2, 3, \dots, F$. Moreover, assume A.2, T.1, $k \geq 2$,¹⁹ and $C''(x) < 0$ for all x . Then, μ^{**} is well defined and $\mu^{**}(\hat{\omega}) = 1$.

PROOF: See Appendix 2.

The intuition underlying Theorem 4.1 is quite apparent from the analysis already conducted in Section 3 for a fixed population of firms. First, recall that when more than one firm exists in the market, the postulated learning dynamics leads them to playing a symmetric MPE. Under increasing returns, this situation can be just temporary since firms make losses. Eventually, some of the incumbents must exit. Since no potential entrant finds entry worthwhile either, this imposes a downward drift on the process until a situation of monopoly is reached. Under these circumstances, the single incumbent eventually learns to play the profit maximizing output. Even though this will recurrently attract other firms into the market, competition among them will make a monopoly situation return relatively fast, thus producing the stated long-run outcome, i.e., most of the time along the process, the market will witness a single firm obtaining monopoly profits. This is to be contrasted with the conclusions obtained under decreasing returns, which are presented in the next subsection.

¹⁸Pursuing the line of proof presented in Appendix 2, it is easy to see that, if the monopoly benefits had several local maxima, one of them will still be selected by a monopoly as the unique long-run outcome of the process. In particular, it would correspond to that one whose basin of attraction contains the output produced in the MPE with two firms.

¹⁹The inequality $k \geq 2$ is required in order to have a meaningful dynamics when there is only one firm in the market.

4.2. Convex Costs

Under decreasing returns, the process will be seen to display an ever-present tendency towards the increase of the firm population size. To have a bound on the number of firms which the market can accommodate, it is standard to contemplate the existence of some fixed costs, no matter how small.

In line with previous notation, let $y^*(n)$ stand for the output produced at the MCPE with n firms. Further denote by $\varpi(x;n)$ the monomorphic state with n firms and output level x . In the present scenario ($C''(\cdot) > 0$), it can be shown that there exists some $n^* \in \mathbb{N}$ such that, provided $C(0)$ is sufficiently small (albeit positive),

$$P(ny^*(n))y^*(n) - C(y^*(n)) \geq 0 \Leftrightarrow n \leq n^*.$$

Thus, n^* represents the maximum number of firms which can co-exist without losses at a symmetric MCPE (here, a Walrasian equilibrium). Denote by $\omega^{**} \equiv \varpi(y^*(n^*); n^*)$ the associated monomorphic state and assume, for simplicity,²⁰ that the profits achieved by firms at it are strictly positive. So, we assume:

T.2 $C(0) > 0$ and $P(n^*y^*(n^*))y^*(n^*) - C(y^*(n^*)) > 0$.

The next result establishes that the state ω^{**} is also the unique long-run state of the process.

THEOREM 4.2: *Assume A.2,²¹ T.2, $n^* \geq 2$, and $C''(x) > 0$ for all x . Then, μ^{**} is well defined and $\mu^{**}(\omega^{**}) = 1$.*

PROOF: See Appendix 2.

The intuition here is polar to the one applicable under increasing returns. On the one hand, for any fixed population of firms, their learning directs them to the Walrasian equilibrium. Under decreasing returns and T.2, whether profits or losses are achieved at this equilibrium depends on the number of firms present in the market. If it is above n^* , population dynamics imposes a downwards adjustment on the existing firms; otherwise, the pressure is upwards and the number of firms increases. Even though the ergodicity of the process guarantees that all configurations are indeed visited with positive frequency (a.s.), those asymmetries among

²⁰This requirement is essentially a convenient (and generic) condition, which could be largely dispensed with in the analysis. Pursuing the line of the proof presented in Appendix 2, it can be shown that, if the Walrasian equilibrium for n^* firms gives exactly zero profits, then the process still selects it unless there exists a sequence of output levels giving strictly negative profits in the corresponding monomorphic states and converging to $y(n^*)$. But, if such a sequence exists, the process will select the Walrasian equilibrium for $n^* - 1$ firms.

²¹Note that A.1 is immediately fulfilled with convex costs.

them induce that, for small ϵ , the process will spend "most of its time" at state ω^{**} .

5. SUMMARY

Our analysis may be summarized as follows.

First, we have found that if the population of firms participating in a market remains fixed (i.e. we do not allow for firm turnover), inter-firm learning dynamics based on imitation and occasional experimentation leads to the following long-run prediction: the market will be largely concentrated in the symmetric Marginal-Cost Pricing Equilibrium. If production returns are decreasing, this outcome represents a Walrasian Equilibrium. However, this is not the case when returns are increasing and the corresponding equilibrium profits are bound to be negative. In a sense, our analysis may be seen as providing a certain learning-based foundation for the usually normative motivation underlying marginal cost pricing.

In a second step in the analysis, we have endogenized the number of firms, allowing for entry and exit to occur in the market in response to the existence of profits or losses. Under certain natural conditions on the relative flexibility of the different decisions, we have found that there is a sharp impact of the technological conditions on the long-run market structure. Specifically, a scenario with decreasing returns displays a drive towards an increasing number of firms and lower profits (obtained at corresponding Walrasian equilibria), while the alternative assumption of increasing returns induces an opposite tendency towards monopoly and high profits. We think that this clear-cut, purely qualitative results may provide some insight into traditional, largely informal, views of how production returns affect market structure.

APPENDIX 1. SUMMARY OF TECHNIQUES

In this section, we will present a summary of the techniques developed by Freidlin and Wentzell (1984), as adapted to our first scenario with a fixed number of firms (cf. the proof of Theorem 3.2). In view of (3)-(5), it should be apparent how to extend them to the framework augmented with entry and exit, as required in the proofs of Theorems 4.1 and 4.2.

Given $\omega = (\omega_1, \dots, \omega_k) \in \Omega$ an ω -tree H is a collection of ordered pairs (or "arrows") (ω^i, ω^j) such that

- i) every $\omega^i \in \Omega \setminus \{\omega\}$ is the first element of one and only one pair
- ii) $\forall \omega^i \in \Omega \setminus \{\omega\}$, there exists a path $\{(\omega^i, \omega^1), (\omega^1, \omega^2), \dots, (\omega^i, \omega)\}$.

The set of all ω -trees is denoted by \mathcal{T}_ω .

Given $\delta > 0$, denote by T_ε the transition matrix of the postulated process with experimentation probability ε . (T_0 will stand for the experimentation-free dynamics.) Given any $\omega \in \Omega$, define:

$$r(\omega) = \sum_{H \in \mathcal{H}_\omega} \prod_{(\omega', \omega'') \in H} T_\varepsilon(\omega', \omega'').$$

Then, provided the process is ergodic (a property which shall be verified in each of our different contexts), the analysis of Freidlin and Wentzell (1984) implies that the (unique) invariant distribution of the process, $\mu \in \Delta(\Omega)$, is given by:

$$\mu(\omega) = \frac{r(\omega)}{\sum_{\omega' \in \Omega} r(\omega')}.$$

As each $r(\omega)$ is a polynomial in ε , it is clear that the limit invariant distribution $\mu_\delta \equiv \lim_{\varepsilon \rightarrow 0} \mu_{\varepsilon, \delta}$ is well defined. To compute the different $r(\omega)$, it is customary in the evolutionary literature (see Kandori, Mailath and Rob (1993) or Young (1993)) to introduce a "cost function" on possible transitions as follows. First, define $d(\omega, \omega')$ as the number of coordinates which differ between ω_k and ω'_k if $T_\varepsilon(\omega, \omega') > 0$. Otherwise, simply make $d(\omega, \omega') = +\infty$. Then, consider²²

$$c : \Omega \times \Omega \rightarrow \mathbb{N}$$

$$c(\omega, \omega') = \min_{\omega'' \in \Omega} \{d(\omega'', \omega') : T_0(\omega, \omega'') > 0\}.$$

Here, $c(\omega, \omega')$ may be interpreted as the minimal number of experimentations needed to take place after a single operation of the experimentation-free dynamics from state ω in order for the process to reach state ω' . The function $c(\cdot)$ may be extended by addition to paths and trees.²³ Then, it is easy to see that the order in ε of the polynomial $r(\omega)$ is given by the minimum number of experimentations required along some ω -tree, i.e. $\min_{H \in \mathcal{H}_\omega} c(H)$. In conclusion, therefore, the stochastically stable states may be singled out as those whose minimum cost trees are themselves minimum across all possible states in Ω .

APPENDIX 2. PROOFS

The statement of Theorem 3.2 in Section 3 concerns the convergence of invariant distributions, i.e. probability measures. In fact, we can prove the following

²²We adhere to the convention that $0 \in \mathbb{N}$.

²³In the framework with a variable population size, transitions involving entry or exit have to incur costs that, in terms of the "units" reflected by the experimentation-based transitions contemplated here, are consistent with conditions (3)-(5).

stronger re-statement, which concerns the convergence of the supports of the invariant distributions.

THEOREM 3.2: *Assume A.1 and A.2. For every $\lambda > 0$ there exists $\delta > 0$ such that if $\delta \leq \delta$,*

$$\text{Support}(\mu_\delta^*) \subset \{\varpi(y) : y \in [(y^* - \lambda, y^* + \lambda) \cap \Gamma(\delta)]\}.$$

PROOF: Denote by \mathcal{A} the collection of recurrent communication classes of the stochastic process T_0 . That is, \mathcal{A} includes all those subsets of Ω which are minimally closed under finite chains of iterations of T_0 (see, for example, Karlin and Taylor (1975)). It is clear that only states belonging to one of these classes can qualify as stochastically stable. Consequently, the characterization of \mathcal{A} provided by the following lemma represents a useful first step in the argument.

LEMMA A.1: *Given δ , the only recurrent communication classes of T_0 are the singletons $\{\varpi(y)\}$, $y \in \Gamma(\delta)$, consisting of monomorphic states.*

PROOF OF LEMMA A.1: Obviously, every monomorphic state defines, as a singleton, a corresponding recurrent communication class of T_0 . To see that no other state can be in a recurrent communication class, it is enough to construct a path which leads from it to some monomorphic state with positive probability.

Denote by $\hat{\Pi}(t) = \max_{j=1, \dots, n} \Pi_j(t)$, i.e., the maximum payoff realized at period t . Moreover, denote by $\Pi^*(t) = \max_{r=0, \dots, k-1} \hat{\Pi}(t-r)$, i.e., the maximum payoff realized within the time range recorded by state $\omega(t)$.

Then, it is claimed that, along any sample path of the process, there is some t such that $\Pi^*(t) = \hat{\Pi}(t) = \Pi_i(t)$ for some $i = 1, 2, \dots, n$. Suppose otherwise, i.e. $\hat{\Pi}(t) < \Pi^*(t)$ for all t . Then, it follows that $\Pi^*(t+k) < \Pi^*(t)$ for all t . Consequently, the sequence $\{\Pi^*(t+l_k)\}_{l=1}^\infty$ is strictly decreasing, which yields a contradiction because the total number of realizable payoffs is finite.

Thus, consider some t such that $\Pi^*(t) = \Pi_i(t)$ and assume w.l.o.g. that $i = 1$. Denote $\hat{x} = x_1(t)$. Suppose now that for the k periods following t , only firm 1 has revision opportunity. This event has positive probability, after which $\omega(t+k)$ will be a repeated state. Further assume that in the next period all firms obtain a revision opportunity and choose exactly the same output level \hat{x} . Again, this event has positive probability. Then, regardless of further revision opportunities, there is also positive probability that $\omega(t+2k) = \varpi(\hat{x})$, which completes the proof of the lemma.

For simplicity in the argument, assume that y^* is an output in $\Gamma(\delta)$ for all δ . (This implies no loss of generality since, otherwise, one of the two "closest" outputs

in the grid may take its place in the analysis, converging to y^* as δ goes to zero.) The remaining part of the proof can be decomposed into the following steps.

- Define $\Phi : [0, K] \times [0, K] \rightarrow \mathbb{R}$ by

$$\begin{aligned}\Phi(x, y) &= [P((n-1)y+x) - C(x)] - [P((n-1)y+x)y - C(y)] \\ &= P((n-1)y+x)(x-y) + (C(y) - C(x))\end{aligned}$$

i.e. the differential profits of a mutant (relative to non-mutants) when deviating from a symmetric situation (y, \dots, y) to a new output level x .²⁴

Note that, obviously, $\Phi(y, y) = 0 \forall y$, and that Φ is a continuously differentiable function with successive partial derivatives:

$$\begin{aligned}\frac{\partial \Phi}{\partial x}(x, y) &= P((n-1)y+x) + P'((n-1)y+x)(x-y) - C'(x) \\ \frac{\partial^2 \Phi}{\partial x^2}(x, y) &= 2P'((n-1)y+x) + P''((n-1)y+x)(x-y) - C''(x).\end{aligned}$$

Consider the function $\Phi_y : [0, K] \rightarrow \mathbb{R}$, given by $\Phi_y(x) = \Phi(x, y)$. If an output y is a local maximum of its own Φ_y , then no close experimentation will destabilize it by achieving better profits. We call such an output *experimentation resistant*. But from the First Order Condition,

$$\Phi'_y(y) = \frac{\partial \Phi}{\partial x}(y, y) = P(ny) - C'(y) = 0 \Leftrightarrow P(ny) = C'(y),$$

that is, for an interior output, the equality of price and marginal cost is a necessary condition for such experimentation resistance. The Second Order Condition is automatically satisfied by A.1:

$$\Phi''_y(y) = \frac{\partial^2 \Phi}{\partial x^2}(y, y) = 2P'(ny) - C''(y) < 0.$$

This means that the only experimentation-resistant (interior) output is y^* . Moreover, define

$$f(y) = \Phi'_y(y) = \frac{\partial \Phi}{\partial x}(y, y) = P(ny) - C'(y)$$

as the slope at y of the differential profit Φ_y . Then, $f(y^*) = 0$ and, provided that $n \geq 2$,

$$\begin{aligned}f'(y) &= nP'(ny) - C''(y) = (n-2)P'(ny) + 2P'(ny) - C''(y) \\ &= (n-2)P'(ny) + \Phi''_y(y)\end{aligned}$$

²⁴Note that if $\Phi(x, y) > 0$, the monomorphic state $\varpi(y)$ may be destabilized by one single mutation to x . For, after one such mutation, there is positive probability that every firm stays with its output for k periods (specifically, if only the mutant receives a revision opportunity during that time span). Then, in the next period, all firms would imitate the mutant if they are able to revise their output.

which is again negative by A.1. Heuristically, this means that close experiments in the direction of y^* are always beneficial. Formally,

$$\begin{aligned}\frac{\partial \Phi}{\partial x}(y, y) &> 0 \Leftrightarrow y < y^* \\ \frac{\partial \Phi}{\partial x}(y, y) &= 0 \Leftrightarrow y = y^* \\ \frac{\partial \Phi}{\partial x}(y, y) &< 0 \Leftrightarrow y > y^*.\end{aligned}$$

- CLAIM 1: $\exists \alpha > 0 : \forall \delta > 0 \wedge \forall x \in (y^* - \alpha, y^* + \alpha) \cap \Gamma(\delta)$, $\exists r \geq 1 : T_r^c(\varpi(x), \varpi(y^*)) > 0$.

PROOF OF CLAIM 1: Let α be such that for every $x \in (y^* - \alpha, y^* + \alpha)$, $x \neq y^*$, $\Phi(x, y^*) < 0$. We know that there exists such an α , because $\Phi(y^*, y^*) = 0$ and y^* is experimentation resistant.

Let $x = y^* + l\delta \in (y^* - \alpha, y^* + \alpha)$ for some $l \in \mathbb{N}$ (the proof is analogous for $x = y^* - l\delta$). We will describe a chain of transitions with positive probability between $\varpi(x)$ and $\varpi(y^*)$.

Let us start this chain with $\varpi(x)$, and suppose that, during l consecutive periods, only one firm is allowed to revise its output, say firm 1, and the rest of the firms mutate downwards. Suppose then that, for k more periods, only firm 1 is allowed to revise but there is no experimentation. Then, since firm 1 has only been able to imitate observed outputs, the state prevailing at the end of this chain will be

$$[(x_1, y^*, \dots, y^*), \dots, (x_k, y^*, \dots, y^*)]$$

with $x_1, \dots, x_k \in [y^*, x] \cap \Gamma(\delta)$. Moreover, since $\Phi(x_i, y^*) < 0$, y^* yields higher profits in each of the configurations (x_i, y^*, \dots, y^*) , $i = 1, \dots, k$.

Therefore in at most s more periods with no experimentation we have that the set $B_1(t+l+s) = \{y^*\}$. Then, there is positive probability that firm 1 receives a revision opportunity, hence choosing y^* . Now, if no firm mutates in the next k periods, the process reaches the monomorphic state $\varpi(y^*)$.

- CLAIM 2: $\forall \lambda > 0, \exists \bar{\delta} > 0 : \forall \delta < \bar{\delta}$,

$$\begin{cases} c(\varpi(x), \varpi(x+\delta)) = 1 & \forall x \in [0, y^* - \lambda] \cap \Gamma(\delta) \\ c(\varpi(x), \varpi(x-\delta)) = 1 & \forall x \in [y^* + \lambda, K] \cap \Gamma(\delta). \end{cases}$$

PROOF OF CLAIM 2: Let

$$D(y) = \begin{cases} \{y' \in (y, K] : \Phi(y', y) > 0 \forall y'' \in (y, y')\} & \text{if } y < y^* \\ \{y' \in [0, y) : \Phi(y', y) > 0 \forall y'' \in (y', y)\} & \text{if } y > y^* \end{cases}$$

and

$$\delta(y) = \begin{cases} | \text{Sup } D(y) - y | & \text{if } y < y^* \\ | \text{Inf } D(y) - y | & \text{if } y > y^*. \end{cases}$$

The claim follows if we prove that

$$\forall \lambda > 0, \exists \delta > 0 : \delta(y) > \delta \forall y \in [0, K] \setminus (y^* - \lambda, y^* + \lambda).$$

To establish this, denote

$$N_{\lambda}^{-} = [0, y^* - \lambda], \quad N_{\lambda}^{+} = [y^* + \lambda, K].$$

By symmetry, it is enough to show the claim for N_{λ}^{-} .

From the analysis of Φ conducted above we already know that, $\forall y \in N_{\lambda}^{-}$, $\Phi(y, y) = 0$ and $\frac{\partial \Phi}{\partial x}(y, y) > 0$, so $D(y) \neq \emptyset$ and $\delta(y)$ is well defined and positive. Note also that, by construction, if $\text{Sup } D(y) < K$, then $\Phi(\text{Sup } D(y), y) = 0$ and $\frac{\partial \Phi}{\partial x}(\text{Sup } D(y), y) \leq 0$.

Suppose now, for the sake of contradiction, that the claim is false. Then, $\forall \tau \in \mathbb{N} \setminus \{0\} \exists y_{\tau} \in N_{\lambda}^{-} : \delta(y_{\tau}) < \frac{1}{\tau}$. Obviously, $\{\delta(y_{\tau})\} \rightarrow 0$. So, for τ big enough, $\text{Sup } D(y_{\tau}) < K$.

As N_{λ}^{-} is compact, $\{y_{\tau}\}$ has a convergent subsequence. Re-indexing, we can just assume, without loss of generality, that $\{y_{\tau}\} \rightarrow \tilde{y} \in N_{\lambda}^{-}$. As $[0, K]$ is also compact, $\{\text{Sup } D(y_{\tau})\}$ has a convergent subsequence and, re-indexing again, we can assume w.l.o.g. that $\{\text{Sup } D(y_{\tau})\} \rightarrow \hat{y} \in [0, K]$. But, since $\{\delta(y_{\tau})\} \rightarrow 0$, it must be that $\tilde{y} = \hat{y}$.

Given that $\frac{\partial \Phi}{\partial x}$ is continuous, $\lim_{\tau \rightarrow \infty} \frac{\partial \Phi}{\partial x}(\text{Sup } D(y_{\tau}), y_{\tau}) = \frac{\partial \Phi}{\partial x}(\tilde{y}, \tilde{y})$. Furthermore, $\frac{\partial \Phi}{\partial x}(\text{Sup } D(y_{\tau}), y_{\tau}) \leq 0$ for τ big enough, which implies that $\lim_{\tau \rightarrow \infty} \frac{\partial \Phi}{\partial x}(\text{Sup } D(y_{\tau}), y_{\tau}) \leq 0$. But $\frac{\partial \Phi}{\partial x}(\tilde{y}, \tilde{y}) > 0$, a contradiction that proves the claim.

- CLAIM 3: $\forall \xi > 0, \forall \lambda > 0, \exists \bar{\delta} > 0 : \forall \delta < \bar{\delta}$,

$$\begin{cases} c(\varpi(x), \varpi(x - \delta)) \geq 2 \forall x \in [\xi, y^* - \xi] \cap \Gamma(\delta) \\ c(\varpi(x), \varpi(x + \delta)) \geq 2 \forall x \in [y^* + \xi, K - \xi] \cap \Gamma(\delta). \end{cases}$$

PROOF OF CLAIM 3: This claim follows from the fact

$$\forall \xi > 0, \lambda > 0, \exists \bar{\delta} > 0 : \forall \delta < \bar{\delta}, \begin{cases} \Phi(x - \delta, x) < 0 \forall x \in [\xi, y^* - \xi] \\ \Phi(x + \delta, x) < 0 \forall x \in [y^* + \xi, K - \xi] \end{cases}$$

and this can be proved by an argument analogous to the one in the proof of Claim 2.²⁵

²⁵The role of ξ is just to avoid boundary problems.

- Now consider any $\lambda > 0$ such that $\lambda < \alpha$, where α is as in Claim 1. Thus, Claim 1 applies to $(y^* - \lambda, y^* + \lambda)$. Now select some $\bar{\delta} > 0$ such that (a) $\bar{\delta} < \frac{\lambda}{2}$, (b) Claim 2 applies for λ , and (c) Claim 3 applies for $\frac{\lambda}{2}$. Then, take any grid $\Gamma(\delta)$ with $\delta < \bar{\delta}$.

By Claim 2, any monomorphic state $\varpi(x)$ with $x \notin (y^* - \lambda, y^* + \lambda)$ can be connected with positive probability to the monomorphic states $\varpi(x')$ with $x' \in (y^* - \lambda, y^* + \lambda)$. But, by Claim 1, the latter can be connected with positive probability to $\varpi(y^*)$. This means that, ultimately, all states can be connected to $\varpi(y^*)$. Therefore, the stochastic process T_t has a unique recurrent communication class, namely the one which contains $\varpi(y^*)$. Obviously all monomorphic states have period one, so this class is aperiodic. Then there exists a unique invariant distribution, which can be characterized using the techniques summarized in Appendix 1.

- By Claim 2, the monomorphic states $\varpi(x)$ with $x \notin (y^* - \lambda, y^* + \lambda)$ can be connected at cost one per state to the next monomorphic state in the direction of $\varpi(y^*)$, until states $\varpi(x')$ with $x' \in (y^* - \lambda, y^* + \lambda)$ are reached. Since it takes at least one experimentation to destabilize a monomorphic state, this connections cannot be obtained at a lesser cost.

Let y_1 and y_2 be the minimum and maximum output levels in the set

$$(y^* - \lambda, y^* + \lambda) \cap \Gamma(\delta),$$

respectively. By Claim 3, $\Phi(y_1 - \delta, y_1) < 0$ and $\Phi(y_2 + \delta, y_2) < 0$, so it takes at least two experimentations to destabilize $\varpi(y_1)$ down or $\varpi(y_2)$ up (but only one to destabilize $\varpi(y_1 - \delta)$ up or $\varpi(y_2 + \delta)$ down).

Let $\Lambda = \{\varpi(y_1), \dots, \varpi(y_2)\}$, and let \hat{H} be a minimal cost tree restricted to the set Λ . Construct now a *complete* tree in Ω , say H , by connecting all the monomorphic states in $\Omega \setminus \hat{H}$ to \hat{H} as explained above, i.e., at cost one per state. Then, connect all non-monomorphic states to the monomorphic ones at cost zero, something which can be done by virtue of Lemma A.1. This will include the arrows $(\varpi(y_1 - \delta), \varpi(y_1))$ and $(\varpi(y_2 + \delta), \varpi(y_2))$.

Again by Lemma A.1, the cost of the tree H just constructed needs to be compared with alternative ω -trees for monomorphic ω only. Consider some such $\varpi(\hat{y}) \notin \Lambda$ (i.e. $|\hat{y} - y^*| > \lambda$). First, note that, in the search for a minimum-cost $\varpi(\hat{y})$ -tree, former considerations allow us to restrict to trees where monomorphic states are directly connected (all other states can be joined to these states at zero cost, in view of Lemma A.1.) In any such $\varpi(\hat{y})$ -tree, every monomorphic state $\varpi(x)$ must be connected either to $\varpi(x + \delta)$ or to $\varpi(x - \delta)$, so that there is only one direction of movement towards the vertex of the tree, $\varpi(\hat{y})$. Taking this into account, it is clear that any

$\varpi(y)$ -tree must have, instead of one of the arrows $(\varpi(y_1 - \delta), \varpi(y_1))$ or $(\varpi(y_2 + \delta), \varpi(y_2))$, the reverse one, effectively increasing the total cost by at least 1 with respect to H . This proves that the stochastically stable states for δ will be in Λ , thus completing the proof of the theorem. $Q.E.D.$

PROOF OF THEOREM 4.1: First, note that in analogy with Lemma A.1, only monomorphic states (now with a variable number of firms present) can be stochastically stable. As explained, this allows us to restrict to ω -trees involving only monomorphic states.

Now let $Mon(n, \delta) \equiv \{\varpi(x; n) : x \in \Gamma(\delta)\}$, where recall that each $\varpi(x; n)$ stands for the monomorphic state with n firms present, all producing output x . Furthermore, given two output levels $x_1, x_2 \in \Gamma(\delta)$, $x_1 < x_2$, denote

$$\begin{aligned} & Branch(x_1, x_2, n, \delta) \\ & \equiv \left\{ \begin{aligned} & (\varpi(x_1; n), \varpi(x_1 + \delta; n)), (\varpi(x_1 + \delta; n), \varpi(x_1 + 2\delta; n)), \dots \\ & \dots, (\varpi(x_2 - 2\delta; n), \varpi(x_2 - \delta; n)), (\varpi(x_2 - \delta; n), \varpi(x_2; n)) \end{aligned} \right\} \end{aligned}$$

as the collection of arrows joining the x_1 to x_2 through all intermediate ones. Symmetrically, we define $Branch(x_2, x_1, n, \delta)$ as the set of converse arrows connecting x_2 to x_1 .

From the proof of Theorem 3.2, it is clear that the minimal cost trees restricted to $Mon(n, \delta)$ are of a very specific form, as summarized by the following claim:

- CLAIM 4: Given $\lambda > 0$, $\exists \delta > 0$: the minimal cost trees in $Mon(n, \delta)$ are of the form

$$E(\widehat{Z}, n, \delta) = Branch(0, y_1, n, \delta) \cup \widehat{Z} \cup Branch(K, y_2, n, \delta),$$

where $y_1 < y_2$ are appropriately chosen outputs (depending on λ , n , and δ) and \widehat{Z} is some minimal cost tree in the set

$$\begin{aligned} A(n, \delta) &= \{\varpi(y; n) : y \in (y^*(n) - \lambda, y^*(n) + \lambda) \cap \Gamma(\delta)\} \\ &= \{\varpi(y; n) : y_1 \leq y \leq y_2, y \in \Gamma(\delta)\}. \end{aligned} \quad (6)$$

Moreover, all arrows in $E(\widehat{Z}, n, \delta) \setminus \widehat{Z}$ have cost exactly equal to 1 while the reverse of any such arrow (except possibly for arrows corresponding to states arbitrarily close to 0 and K)²⁶ would have a cost of at least 2.

²⁶As in the proof of Theorem 3.2, these states will not interfere with the essential part of the argument since they can be restricted to lie arbitrarily close to the extremes of the interval – recall Claim 3.

Under increasing returns to scale, the states $\varpi(y^*(n); n)$ yield losses for all $n \geq 2$. By continuity, therefore, we can choose λ above small enough so that for all $n \geq 2$ and every $y \in (y^*(n) - \lambda, y^*(n) + \lambda)$, the monomorphic states $\varpi(y; n)$ also induce losses.

On the other hand, the state $\hat{\omega}$ yields (positive) profits. Thus, again by continuity, $\exists B > 0$ such that $\forall y \in (\hat{y} - B, \hat{y} + B)$, $\varpi(y; 1)$ induces profits. Furthermore, given this B , suppose δ above is chosen small enough so that $\frac{B}{\delta} \geq q + 1$, where q is as in (4).

The proof is now straightforward. First, we construct an $\hat{\omega}$ -tree which we will show has minimal cost. To do so, consider, $\forall n \geq 2$, minimal cost trees $E(\widehat{Z}, n, \delta)$ restricted to $Mon(n, \delta)$. For $n = 1$, it is obvious from Assumption T.1 that the minimal cost tree restricted to $Mon(1, \delta)$ is

$$E(1, \delta) = Branch(0, \hat{y}, 1, \delta) \cup Branch(K, \hat{y}, 1, \delta).$$

Note that the requirement $\frac{B}{\delta} \geq q + 1$ means that, if

$$(\hat{y} - B, \hat{y} + B) \cap \Gamma(\delta) = \{y'_1, y'_1 + \delta, \dots, y'_2 - \delta, y'_2\},$$

then $Branch(y'_1, \hat{y}, 1, \delta)$ and $Branch(y'_2, \hat{y}, 1, \delta)$ each have at least $q + 1$ arrows. Thus, reversing all the arrows in one of such a branch will increase the cost by at least $q + 1$ which, by (5), is greater than the cost associated to single entry or exit of firms.

Second, we join all the trees $E(\widehat{Z}, n, \delta)$ thus constructed and the tree $E(1, \delta)$ as follows:

- $\forall n \geq 2$, join the vertex of $E(\widehat{Z}, n, \delta)$ to a state with $n - 1$ firms through a single exit event, whose probability σ is an infinitesimal of order no smaller than ϵ by (3). This is possible because the vertex of each $E(\widehat{Z}, n, \delta)$ yields losses, as observed above.
- Connect the state $\varpi(\theta_0)$ involving no firms to a state with one firm through a single entry event, whose probability η is an infinitesimal of order larger than ϵ by (3) and (5).

In this manner, we have constructed a full $\hat{\omega}$ -tree, say \hat{H} , which involves (i) a minimal cost at each “level” $Mon(n, \delta)$; (ii) one exit per level; and (iii) one entry from $\varpi(\theta_0)$. We now claim that, given any other monomorphic state $\tilde{\omega}$, it is not possible to construct an $\tilde{\omega}$ -tree with the same or lower cost than \hat{H} . Consider any other state $\tilde{\omega}$.

- If $\tilde{\omega}$ involves $\tilde{n} \geq 2$ firms, then any $\tilde{\omega}$ -tree will include at least one entry of a firm from a state in $Mon(1, \delta)$ instead of an exit from a state in $Mon(2, \delta)$. But, as exit is an event whose probability is an infinitesimal of lower order than entry (see (5)), this results in a cost strictly greater than the cost of \hat{H} . (Recall that the cost realized within each level $Mon(n, \delta)$ has to be no

lower than in \hat{H} by construction, and that all "levels" $Mon(n, \delta)$ with $n > \tilde{n}$ require at least one single entry or exit to be joined to a full tree.)

- If $\tilde{\omega}$ involves $\tilde{n} = 1$ firms, then any $\tilde{\omega}$ -tree will have a cost greater than \hat{H} by Assumption T.1.

- If $\tilde{\omega} = \varpi(\theta_0)$, then any $\tilde{\omega}$ -tree will include an exit from a state in $Mon(1, \delta)$ instead of an entry from $\varpi(\theta_0)$, which indeed saves some cost as compared to that of \hat{H} (recall (5)). However, since exit is only possible from a state with losses, such $\tilde{\omega}$ -tree may be completed in only two ways:

One possibility is to rely on a sub-tree within $Mon(1, \delta)$ whose vertex state displays losses and thus may be connected to $\varpi(\theta_0)$. This will involve reversing all the arrows in either $Branch(y'_1, \hat{y}, 1, \delta)$ or $Branch(y'_2, \hat{y}, 1, \delta)$ which, as observed above, is more costly than one single entry or exit. Therefore, it results in a cost greater than the cost of \hat{H} .

The second possibility involves connecting a state in $Mon(1, \delta)$ to another state in some $Mon(n, \delta)$, $n > 1$, also connecting a state in some $Mon(n', \delta)$, $n' \geq 1$, to either $\varpi(\theta_0)$ or another state in some $Mon(n'', \delta)$ with $1 \leq n'' < n'$. In view of (5), such operations will amount to a cost strictly greater than that of \hat{H} .

This proves that $\tilde{\omega}$ is the only stochastically stable state, thus completing the proof of the theorem. Q.E.D.

PROOF OF THEOREM 4.2: As for the proof of the former results (recall Lemma A.1), we may restrict our application of the graph-theoretic techniques to monomorphic states, which are the only candidates to being stochastically stable. Furthermore, relying on the notation introduced in the proof of Theorem 4.1, we also note that Claim 4 above applies without modification to the present context.

By Assumption T.2, the states $\varpi(y^*(n); n)$ all give losses for $n > n^*$, and profits for $n \leq n^*$. Then, the two following statements are true by continuity:

- (1) $\exists B_1 > 0 : \forall n > n^*, \forall y \in (y^*(n) - B_1, y^*(n) + B_1)$, $\varpi(y; n)$ yields losses;
- (2) $\exists B_2 > 0 : \forall n \leq n^*, \forall y \in (y^*(n) - B_2, y^*(n) + B_2)$, $\varpi(y; n)$ yields profits.

Choose now $0 < \lambda < B = \min\{B_1, B_2\}$ and find $\delta > 0$ such that Claim 4 in the proof of Theorem 4.1 holds $\forall n \geq 2$ and $\frac{B-\lambda}{\delta} > q + 3$, where q is as in (4).

The minimal cost trees $E(\tilde{Z}, n, \delta)$ can be decomposed as follows:

$$E(\tilde{Z}, n, \delta) = Branch(0, y'_1, n, \delta) \cup Branch(y'_1, y_1, n, \delta) \cup \tilde{Z} \cup \\ \cup Branch(y'_2, y_2, n, \delta) \cup Branch(K, y'_2, n, \delta)$$

with \tilde{Z} , y_1 , y_2 as in Claim 4 and

$$(y^*(n) - B, y^*(n) + B) \cap \Gamma(\delta) = \{y'_1, y'_1 + \delta, \dots, y'_2 - \delta, y'_2\}.$$

Note that, given $n \geq 2$, the requirement $\frac{B-\lambda}{\delta} > q + 3$ means that there are at least $q + 1$ arrows in $Branch(y'_1, y_1, n, \delta)$ and $Branch(y'_2, y_2, n, \delta)$. Thus, reversing all of the arrows in one of these branches will increase the cost by at least $q + 1$ which, by (4), is greater than the cost associated to a single entry or exit by firms.

Let $\tilde{\omega}$ denote the vertex of some minimal-cost tree $E(\tilde{Z}, n^*, \delta)$, as introduced above. The remainder of the proof now relies on the construction of an $\tilde{\omega}$ -tree with minimal cost, which is compared with any other alternative tree. This construction involves the following steps.

First, $\forall n \geq 2$, $n \neq n^*$, identify minimal cost trees $E(\tilde{Z}, n, \delta)$ restricted to $Mon(n, \delta)$, as indicated above.

Second, choose some minimal cost tree in $Mon(1, \delta)$, say $E(1, \delta)$, with its vertex state displaying (positive) profits.

Third, join all the previous trees $E(\tilde{Z}, n, \delta)$ and $E(1, \delta)$ as follows:

- $\forall n > n^*$, connect the vertex of $E(\tilde{Z}, n, \delta)$ to a state with $n - 1$ firms through a single exit event, whose probability σ is an infinitesimal of order no smaller than ϵ by (3). This is possible because the vertex state in such $E(\tilde{Z}, n, \delta)$ yields losses, as observed above.
- $\forall n < n^*$, connect the vertex of $E(\tilde{Z}, n, \delta)$ to a state with $n + 1$ firms through a single entry event, whose probability η is an infinitesimal of order larger than ϵ by (3) and (5). This is possible because the vertex state in such $E(\tilde{Z}, n, \delta)$ induces profits, as observed above as well.

(c) Connect the state $\varpi(\theta_0)$ to a state with one firm with a single entry event (with probability η).

Through (a)-(c) we construct a full $\tilde{\omega}^{**}$ -tree, say \hat{H} , that involves: (i) a minimal cost at each "level" $Mon(n, \delta)$; (ii) one exit per level above n^* ; (iii) one entry per level for $n < n^*$; (iv) one entry from $\varpi(\theta_0)$.

We now claim that, given any other state $\tilde{\omega}$, it is not possible to construct an $\tilde{\omega}$ -tree with the same or lower cost than \hat{H} . Consider any other state $\tilde{\omega}$.

- If $\tilde{\omega}$ involves $\tilde{n} > n^*$ firms, then any $\tilde{\omega}$ -tree will include at least one entry of a firm from a state in $Mon(n^*, \delta)$ instead of an exit from a state in $Mon(n^* + 1, \delta)$. As exit is measured by an infinitesimal of lower order than entry (see (5)), this induces a cost strictly greater than that of \hat{H} . (Recall that the cost realized within each level $Mon(n, \delta)$ has to be no lower than that in \hat{H} by construction, and that all levels $Mon(n, \delta)$ with $n > \tilde{n}$ require at least one single entry or exit to be joined to a full tree.)

- If $\tilde{\omega}$ has $\tilde{n} < n^*$ firms, then any $\tilde{\omega}$ -tree will involve an exit from a state in $Mon(n^*, \delta)$ instead of an entry from a state in $Mon(n^* - 1, \delta)$, which saves some cost as compared with \hat{H} (recall (5)). However, as exit occurs only from states which yield losses, such $\tilde{\omega}$ -tree can be constructed only in two ways.

The first possibility is to rely on a sub-tree within $Mon(n^*, \delta)$ whose vertex state yields losses and thus this vertex state may be connected to some state in $Mon(n^* - 1, \delta)$. This will involve reversing all the arrows in either $Branch(y'_1, y_1, n, \delta)$ or $Branch(y'_2, y_2, n, \delta)$ which, as observed above, is more costly than one single entry or exit. Therefore, it results in a cost greater than that of \tilde{H} .

The second possibility involves connecting a state in $Mon(n^*, \delta)$ to another state in some $Mon(n, \delta)$, $n > n^*$, also connecting a state in some $Mon(n', \delta)$, $n' \leq n$, to another state in some $Mon(n'', \delta)$ with $n'' < n^*$. In view of (3)-(5), such operations will amount to a cost strictly greater than that of \tilde{H} .

- If $\tilde{\omega}$ has n^* firms and is a vertex of a minimal-cost tree in $Mon(n^*, \delta)$, it must belong by construction to the set $A(n^*, \delta)$, as defined in (6). Then, $\tilde{\omega}$ may equivalently fulfill the role of $\tilde{\omega}$ in the final part of the proof (see below), thus leading as well to the desired conclusion. Otherwise, i.e. if it is not a vertex state of any minimal-cost tree in $Mon(n^*, \delta)$, it immediately follows that any $\tilde{\omega}$ -tree must involve a cost greater than \tilde{H} .

• Finally, if $\tilde{\omega} = \varpi(\theta_0)$, then to construct an $\tilde{\omega}$ -tree it would be necessary either to include an exit from a state in some $Mon(n, \delta)$, $n \leq n^*$, or consider an entry from a state in some $Mon(n, \delta)$, $n \geq n^*$, and an exit from a state in some $Mon(n', \delta)$, $n' \geq n$. In either case, the resulting cost must be greater than that of \tilde{H} , in view of (3)-(5).

From the above discussion it follows that every stochastically stable state must belong to the set $A(n^*, \delta)$, as defined in (6) for any given λ . By choosing $\lambda \rightarrow 0$ and, correspondingly, $\delta \rightarrow 0$, we have $A(n^*, \delta) \rightarrow \{\omega^{**}\}$, which completes the proof. *Q.E.D.*

REFERENCES

- ALCHIAN, A. (1950): "Uncertainty, Evolution, and Economic theory," *Journal of Political Economy*, 58, 211-221.
- BAUMOL, W., J. PANZAR, AND R. WILLIG (1982): *Contestable Markets and the Theory of Industry Structure*. New York: Harcourt Brace Jovanovich.
- BJÖRNERSTEDT, J., AND K. SCHLAG (1996): "On the Evolution of Imitative Behavior," University of Bonn, Discussion Paper No. B-378.
- FREIDLIN, M.I., AND A.D. WENTZEL (1984): *Random Perturbations of Dynamical Systems*. New York: Springer Verlag.
- FRIEDMAN, J. (1982): "Oligopoly Theory," in *Handbook of Mathematical Economics*, Vol. II, ed. by K.J. Arrow and M.D. Intriligator. Amsterdam: North-Holland.

- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutations and Long-Run Equilibria in Games," *Econometrica*, 61, 29-56.
- KARLIN, S., AND H.M. TAYLOR (1975): *A First course in Stochastic Processes*. London: Academic Press.
- NOVSHAK, W. (1980): "Cournot Equilibrium with Free Entry," *Review of Economic Studies*, 47, 473-486.
- QUINZIN, M. (1992): *Increasing Returns and Efficiency*. New York: Oxford University Press.
- RHODE, P., AND M. STEGEMAN (1994): "Non-Nash Equilibria of Darwinian Dynamics (with applications to duopoly)," mimeo, University of North Carolina.
- ROBSON, A., AND F. VEGA-REDONDO (1996): "Efficient Equilibrium Selection in Evolutionary Games with Random Matching," *Journal of Economic Theory*, 70, 65-92.
- SAMUELSON, L. (1997): *Evolutionary Games and Equilibrium Selection*. Cambridge: MIT Press.
- SOHAFER, M. (1989): "Are Profit Maximizers the Best Survivors? A Darwinian Model of Economic Natural Selection," *Journal of Economic Behavior and Organization*, 12, 29-45.
- SCHLAG, K. (1998): "Why Imitate and If So, How? A Bounded Rational Approach to Multi-armed Bandits," *Journal of Economic Theory*, 78, 130-156.
- VEGA-REDONDO, F. (1996): *Evolution, Games, and Economic Behaviour*. Oxford: Oxford University Press.
- (1997): "The Evolution of Walrasian Behavior," *Econometrica*, 65, 375-384.
- YOUNG, P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57-84.
- C. Alos-Ferrer, Department of Economics, University of Vienna, Hohenstaufengasse 9, A-1010 Vienna, Austria; e-mail: Carlos.Alos-Ferrer@univie.ac.at
- A.B. Ania, Department of Economics, University of Vienna, Hohenstaufengasse 9, A-1010 Vienna, Austria; e-mail: Ana-Begona.Ania-Martinez@univie.ac.at
- F. Vega-Redondo, Facultad de Económicas, University of Alicante and Instituto Valenciano de Investigaciones Económicas, 03071 Alicante, Spain; e-mail: fvega@merlin.fae.ua.es